

**Polygon Triangulation in $O(N \text{ Log Log } N)$ Time
With Simple Data Structures**

by

David G. Kirkpatrick, University of British Columbia

Maria M. Klawe, University of British Columbia

and

Robert E. Tarjan, Princeton University and

NEC Research Institute

TR 90-17

Technical Report 90-17

June, 1990

Abstract

We give a new $O(n \log \log n)$ -time deterministic algorithm for triangulating simple n -vertex polygons, which avoids the use of complicated data structures. In addition, for polygons whose vertices have integer coordinates of polynomially bounded size, the algorithm can be modified to run in $O(n \log^* n)$ time. The major new techniques employed are the efficient location of horizontal visibility edges that partition the interior of the polygon into regions of approximately equal size, and a linear-time algorithm for obtaining the horizontal visibility partition of a subchain of a polygonal chain, from the horizontal visibility partition of the entire chain. The latter technique has other interesting applications, including a linear-time algorithm to convert a Steiner triangulation of a polygon into a true triangulation.

This research was partially supported by the following grants: NSERC 583584, NSERC 580485, ONRN00014-87-0467, and by DIMACS, an NSF Science and Technology Center (NSF-STC88-09648).

I. Introduction

A **triangulation** T of a simple polygon P with n edges is a partition of its interior into exactly $n - 2$ triangles. We refer to the edges of $T \setminus P$ as **chords**. The problem of triangulating a simple polygon; that is, determining a set of exactly $n - 3$ non-intersecting chords, has attracted considerable attention in the computational geometry literature and elsewhere. Determining the inherent complexity of this problem remains one of the outstanding open problems in this field.

The practical motivation for developing efficient triangulation algorithms is based on the widespread applications of this problem in related disciplines such as computer graphics [FM84, FFR83, L81], pattern recognition [T80] and computational morphology [T86, T88₂] as well as its role as a fundamental building block in the solution of a variety of (superficially) more complex problems in computational geometry such as geometric decomposition [FM84, LTL89], visibility [C82, GHLST87], shortest path [GHLST87], separability [BT86], and subdivision preconditioning [EGS86, K83] problems.

An $O(n \log n)$ upper bound on the complexity of triangulating arbitrary simple polygons was first established by Garey et al. [GJPT78]. Their algorithm is based on a linear-time method for triangulating monotone polygons. Like the subsequent triangulation algorithm of Chazelle [C82], it involves an explicit sorting step. The first improvements on the $O(n \log n)$ bound came by expressing the complexity in terms of structural characteristics of the input polygon (such as the number of reflex angles [HM83, FM84] or its sinuosity [CI84]) or the output [T88]. A separation between the worst-case complexities of polygon triangulation and sorting was not established, however, until Tarjan and van Wyk [TV88] presented an $O(n \log \log n)$ algorithm for the former. Subsequently, Clarkson et al [CTV88] described a randomized algorithm that gives an $O(n \log^* n)$ upper bound on the expected time complexity of the problem.

Other notable developments expanded the class of families of simple polygons for which linear-time algorithms were known [TA82], established a family of problems which are linear-time equivalent to polygon triangulation [FM84], and specified some fundamental primitives out of which efficient triangulation algorithms could be constructed [C82, CI84, HMRT86, FNTV88].

Very recently [C90], Chazelle has given an $O(n)$ deterministic algorithm, thus finally settling the question of the asymptotic complexity of polygon triangulation. Chazelle's result is very beautiful and unquestionably will stand as one of the landmarks of computational geometry for many years to come. Although the underlying ideas in Chazelle's algorithm are both simple and elegant, at least in its current state the algorithm uses some fairly substantial machinery, making an intuitive understanding of its correctness a fair challenge for the non-expert.

Our contributions in this paper are best appreciated in conjunction with these related results. We present a new $O(n \log \log n)$ -time algorithm for triangulating simple polygons that uses only elementary data structures and is formulated in terms of three basic geometric subproblems (two of which, including planar subdivision search, have been identified and studied previously), each of which is of interest in its own right.

Like all of the recent approaches to triangulation [CI84, FM84, TV88, CTV88, C90], ours focuses on the problem of constructing the horizontal visibility partition (HVP) of the edges of a given polygonal chain, i.e. the partition of the plane obtained by adding horizontal edges connecting each vertex to the closest point on the chain on each side if such a point exists, and to a point at infinity otherwise. This problem is well known to be linear-time equivalent to the problem of triangulating a simple polygon [FM84, CI84].

Two natural operations on HVPs form the building blocks of our algorithm. First we need to be able to construct the HVP of a polygonal chain P given the HVP of some prefix subchain P_1 of P together with the HVP of the remainder $P \setminus P_1$. We refer to this as a **merging** of HVPs. Chazelle and Incerpi [CI84] show that merging of HVPs can be done in time linear in the size of the input partitions. They use this as the merge step in a natural divide-and-conquer algorithm for building the HVP of an arbitrary chain. We use merging in a fundamentally different way. Our algorithm uses a balanced divide-and-conquer approach somewhat like that of [TV88], but our splitting is achieved without Jordan sorting and without the use of finger search trees, which are (both conceptually and practically) the most demanding components of the earlier algorithm.

In addition to the merge operation on HVPs we also need to perform its inverse, which we call a **splitting** operation. Specifically given an HVP for P we need to construct the HVP for P_1 , a prefix of P . We show that this can also be done in linear time.

The splitting operation can be viewed as the problem of updating an HVP upon the removal of certain of its non-horizontal edges. Exploiting the linear-time equivalence between polygon triangulation and HVP construction, our splitting algorithm has the following interesting implication for the dynamic maintenance of arbitrary point-set triangulations: given a triangulation T of a point set S and a subset $D \subseteq S$ such that the connected components of the set formed by taking the union of the triangles (including their interiors) adjacent to the vertices in D are simply connected, it is possible to update T to a triangulation of $S \setminus D$ in time $O(\text{deg}_T(D))$ where $\text{deg}_T(D)$ denotes the sum, over all vertices v in D , of the degree of v in T . This immediately implies that the problem of triangulating a simple polygon **allowing Steiner points** is linear-time equivalent to the (Steiner-point-free) polygon triangulation problem.

The remainder of the paper is organized as follows. Section 2 gives the $O(n \log \log n)$ algorithm for computing the horizontal visibility partition of a simple polygon with n vertices. Section 3 describes the linear-time HVP splitting algorithm. In Section 4 we show how the algorithm can be modified to run in $O(n \log^* n)$ time for polygons whose vertices have integer coordinates of polynomially bounded size.

2. The Horizontal Visibility Partition Algorithm

In this section we give our $O(n \log \log n)$ algorithm for computing the horizontal visibility partition of a simple polygon with n vertices. We work with a variant of the horizontal visibility partition called the wrap-around partition, and we compute the wrap-around partitions of directed polygonal chains as well as polygons. Moreover, we assume that the chains and polygons have no self-crossings, and are non-degenerate in the sense that two distinct vertices cannot have the same vertical coordinate. None of these modifications alters the asymptotic complexity of the problem.

If e is a directed edge, the **positive side** of e is its right-hand side with respect to its direction. The positive side of a directed chain is defined analogously. The horizontal line through a vertex on a directed chain C may intersect C on its positive sides from 0, 1 or 2 directions, as is illustrated in Figure 1.




Figure 1.

Wrap-around partitions regard chains and polygons as being embedded on an infinite vertical cylinder, in which horizontal lines going off to infinity wrap-around and reappear on the other side. The **one-sided wrap-around partition** of a directed polygonal chain C is the partition of the cylinder obtained by adding for each vertex v of C , each horizontal line (if any) which starts at v and continues on a positive side of C at v , until it intersects C again. Note that if C is the boundary of a polygon directed so its positive side is the polygon's interior side, then the one-sided wrap-around partition is the usual horizontal visibility partition of the polygon. The **doubling** of a directed chain is the simple polygon obtained by taking the concatenation of C with a slight perturbation of its reversal, with the perturbation chosen so that the doubling satisfies the non-degeneracy condition and so that the perturbed reversal lies on the negative side of C . The **two-sided wrap-around partition** adds both sets of horizontal edges from the one-sided partitions of C and its (unperturbed) reversal. Although our construction could be modified to deal only with one-sided wrap-around partitions, a number of the steps seem both simpler and more intuitive in the context of two-sided partitions. For this reason we use the "doubling" trick to move between one-sided and two-sided partitions. For any point x on a chain C , its **horizontal neighbours** are the points on the intersection of C with the horizontal line through x which are closest to x on each side. The interior horizontal neighbours of a point on a polygon are defined analogously. These definitions are illustrated in Figure 2. For convenience, we will abbreviate wrap-around partition by **WP**.

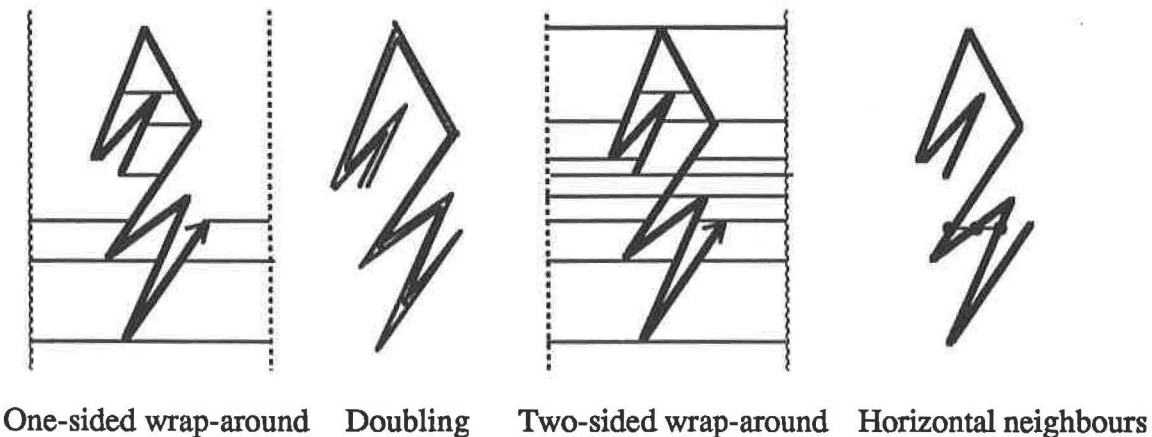


Figure 2.

We will use the term **length** of a chain or polygon to mean the number of its edges. A **k -uniform partition** of a chain or polygon is a partition (with all divisions occurring at vertices) into segments of length between $k/2$ and k inclusive, such that we know the two-sided WP of each segment in the partition. Note that the number of segments in a k -uniform partition of an r -vertex chain or polygon is between r/k and $2r/k$. The main component of our algorithm is the following theorem.

Theorem 2.1. Suppose $k \geq (r/2)^{2/3}$. Given a k -uniform partition of an r -vertex non-degenerate polygon P , we can compute the WP of the interior of P in $O(r)$ time.

Given Theorem 2.1, it is easy to give an $O(n \log \log n)$ -time algorithm for computing the two-sided WP of an n -vertex non-degenerate chain. We partition C into segments of length between $n^{2/3}/2$ and $n^{2/3}$, and use our algorithm recursively to compute the two-sided WP of each of the segments. Next, by Theorem 2.1, regarding the doubling of C as the boundary of a $2n$ -vertex polygon, P , in $O(n)$ time we compute the WP of the interior of P , from which we can easily obtain the two-sided WP of C in $O(n)$ time by correcting the perturbations made in forming the doubling of C . The time, $T(n)$, required by this algorithm satisfies the recurrence $T(n) = O(n) + \sum_i T(n_i)$ where $\sum_i n_i = n$ and $n_i \leq n^{2/3}$, which implies $T(n) = O(n \log \log n)$. The choice of $2/3$ as the exponent is arbitrary, and could have been any number strictly between $1/2$ and 1 .

Proof of Theorem 2.1.

The overall idea is as follows. The theorem is proved by induction on r . We fix a constant r_0 , specified in Lemma 2.7, and for $r < r_0$ compute the WP of P by merging together all the WPs of the segments in the uniform partition. Thus we assume $r \geq r_0$. We wish to partition the interior of P into regions by adding horizontal edges joining some of the vertices on P to their interior horizontal neighbours. We will call these regions **chunks**, and will compute the WP of P by computing the WP of each chunk separately. To achieve the $O(r)$ running time, we want the partition into chunks to have the property that at least a fixed fraction of the vertices of P will lie on regions whose WP we can compute easily from the WP information about the segments. To do this we introduce the notion of special point. A point x is a **special point** of a segment in a k -uniform partition of P if it has the following two properties:

- (i) x has an interior horizontal neighbour $h(x)$ in P such that if P is cut at x and $h(x)$, both the resulting subpaths will contain at least $k/36$ of the vertices of P . We will call $h(x)$ the special horizontal neighbour of x (note that x might have more than one horizontal neighbour).
- (ii) if the segment is cut at x both the resulting subsegments will contain (counting x if x is a vertex) at least $k/36$ of the vertices of the segment.

The complete proof of the theorem depends on several straightforward facts about special points, k -uniform partitions and chunks, which we prove later in this section. In particular, by proving that each segment contains a special point (Lemma 2.6), that we can identify a special point in a segment in $O(k)$ time (Lemma 2.6), and that we can find the special horizontal neighbours in P of all the special points in $O(r)$ total time (Lemma 2.2), we see that in $O(r)$ time we can partition the interior of P into chunks by adding the horizontal edges joining each special point to its special horizontal neighbour.

Figure 3 illustrates the segments, special points and chunks as they are used in the algorithm.




Figure 3.

The boundary of a chunk alternates between horizontal edges and pieces of the boundary of P , where each such piece is the concatenation of at most two subsegments of the segments in the uniform partition. We will call the number of horizontal edges on the boundary the **degree** of a chunk. We will call a chunk **good** if it has degree at most 2, and **bad** otherwise. Using the linear-time splitting and merging algorithms, we can compute the WP of a good chunk in $O(k)$ time, and hence we compute the WPs of all the good chunks in $O(r)$ time. Using the linear-time splitting and merging algorithms, we will show that in $O(r)$ time we can obtain a k -uniform partition of the boundary of each bad chunk. The bad chunks are degenerate since any pair of vertices forming the endpoints of an added horizontal edge obviously have the same vertical coordinate, but the bad chunks can easily be made non-degenerate by slightly perturbing one endpoint of each horizontal edge. After doing this, we then complete the computation of the WP of P by applying Theorem 2.1 inductively to each bad chunk.

Finally, in order to prove that the total running time used in the computation of the WP of P is $O(r)$, it suffices to show that the total number of vertices lying on the boundaries of bad chunks is less than cr for some fixed constant $c < 1$, which we do in Lemma 2.7. ■

We now prove the lemmas needed for Theorem 2.1.

Lemma 2.2. Given any set of $O(r/k)$ points x_1, \dots, x_s on the boundary of P we can find their interior horizontal neighbours in $O(r)$ time.

Proof. Let S_i be the segments in the uniform partition, and let $WP(S_i)$ denote the WP of S_i . By using planar point location [EGS86], in $O(r)$ time we create a data structure for each $WP(S_i)$, so that for any point x and each i , in $O(\log |S_i|)$ time we can locate the region of $WP(S_i)$ containing x . We next find, in $O((r/k)^2 \log k) = O(r)$ time (since $k \geq (r/2)^{2/3}$), the interior horizontal neighbours in P of each x_i , as follows. For each x_i and each S_j we find the region in $WP(S_j)$ containing x_i and the two edges, $L_j(x_i)$ and $R_j(x_i)$ which x_i sees to its left and right, respectively, in the region. Using the planar point location data structures this can be done in $O((r/k)^2 \log k)$ time in total. Now for each x_i , depending on where the interior of P is with respect to x_i , we examine one or both of the two sets of edges $\{L_j(x_i)\}$ and $\{R_j(x_i)\}$ to find the edge(s) which x_i sees in the interior of P . This requires at most $O((r/k)^2) = O(r)$ time. ■

Lemma 2.3. If a polygon C is partitioned into t segments of length at most k whose two-sided WPs we know, then in $O(tk)$ time we can produce a k -uniform partition of C .

Proof. Let $x_1 \dots, x_s$ be vertices in clockwise order on C that partition C so that all segments have length between $k/2$ and k . We will call these the uniform segments. Note that $s/2 \leq t$. We will show that we can find the two-sided WPs of the uniform segments in $O(tk)$ time. We first split into subsegments the original segments whose two-sided WPs we know, by splitting at each x_i and computing the WP of each of these subsegments using the linear-time splitting algorithm. Since each segment that gets split has length at most k and s splits are performed, the time used in splitting is $O(sk) = O(tk)$. We compute the WP of each uniform segment by merging the WPs of its subsegments. Since each merge involves chains of length at most k and at most t merges are performed, the total time used is $O(tk)$. ■

Lemma 2.4. The total time needed to produce a k -uniform partition of the boundary of every chunk is $O(r)$.

Proof. Since the chunks are formed by adding at most $2r/k$ horizontal edges, in $O(r)$ time we can split each segment in the k -uniform partition of P at the endpoints of the added horizontal edges and compute the WPs of these subsegments. Note that the total number of subsegments is still $O(r/k)$ and that every subsegment is of length at most k . Now, using Lemma 2.3, the total time to produce a k -uniform partition of the boundary of every chunk is $O(r)$. ■

Before proving the remaining facts about special points and chunks we need to discuss the dual graphs associated with visibility partitions. Given a WP of a chain (either one-sided or two-sided) we define the **dual graph of the partition** to be the graph whose vertices are the connected components of the partition of the cylinder induced by the WP, such that two vertices are adjacent if their components share a horizontal edge of the partition. Figure 4 shows examples. The main reason we use wrap-around partitions is the easily proved fact that the dual of an (either one-sided or two-sided) WP of a chain is always a tree. In addition, the non-degeneracy condition implies that the tree has maximum degree at most 4. The dual graph of a one-sided WP of a polygon is defined similarly except that the vertices are restricted to the connected components lying in the interior of polygon, and again the dual graph is always a tree of maximum degree at most 4.

More generally, the dual graphs of partitions obtained by adding any set of edges joining points to their (wrap-around) horizontal neighbours are trees, though not necessarily of maximum degree four. Note that the degree of a chunk is simply the degree of the corresponding vertex in the dual graph of the chunk partition.




Figure 4.

Lemma 2.6. If N is a segment in a k -uniform partition of a polygon P , then N has a special point, and we can identify it in $O(k)$ time.

Proof. Let Q be the middle subsegment of N obtained by chopping off $\lfloor k/36 \rfloor$ edges at each end of N . The length of Q is at least $\lfloor 4k/9 \rfloor$ since the length of N is at least $k/2$. It is easy to see that every point of Q satisfies special point property (ii), so it suffices to show that some point of Q satisfies special point property (i) and that given the WP of N we can identify such a point in $O(k)$ time.

Let T be the set of edges in the dual graph of the WP of the interior of P whose corresponding horizontal edges have at least one endpoint lying on Q . Since Q is a connected portion of the boundary of P , T is connected and hence is a tree of degree at most 4. It is not hard to prove that every pair of consecutive edges of Q must contain an endpoint of a horizontal edge of the WP of P and hence there are at least $\lfloor 2k/9 \rfloor$ edges in T . Now, since the maximum degree of T is 4, it is easy to prove that there is some edge e in T such that both connected components in $T - e$ have at least $k/18 - 1$ edges. Moreover, given T we can find e in time linear in the size of T . Let T_1, T_2 be the two connected components of $T - e$, let x be an endpoint of the horizontal edge corresponding to e which lies on Q , and let $h(x)$ be the other endpoint. If we cut P at x and $h(x)$, then the resulting subpaths of P correspond in an obvious way to T_1 and T_2 . The number of vertices on a subpath is at least the number of vertices which are an endpoint of a horizontal edge corresponding to an edge in $\{e\} \cup T_i$, where T_i is the appropriate connected component of $T - e$. Since each vertex is the endpoint of at most two horizontal edges corresponding to edges in T , we see that each subpath has at least $k/36$ vertices, and hence x satisfies property (i).

Finally, we must show that we can identify x in $O(k)$ time. Let P' be the polygon obtained by doubling N . Applying exactly the same argument regarding Q as a subsegment of P' shows that in $O(k)$ time we can find a point x on Q with an interior horizontal neighbour $h'(x)$ on P' such that cutting P' at x and $h'(x)$ yields two subpaths with at least $k/36$ vertices on each subpath. We claim that x is a special point of N . As before, since x is on Q it suffices to show x satisfies property (i). Let $h(x)$ be the interior horizontal neighbour of x on P in the same direction as $h'(x)$. If $h(x) = h'(x)$ we are done. If not, then $h(x)$ is not on

N . Now since x is at least $\lfloor k/36 \rfloor$ edges away from both ends of N , it is clear that cutting P at both x and $h(x)$ results in subpaths of P with at least $k/36$ vertices. ■

Lemma 2.7. There are constants $c < 1$ and r_0 such that if $r \geq r_0$ then at most cr vertices lie on the boundaries of bad chunks.

Proof. The total number of vertices lying on all chunks is at most $r + 8r/k$ since at most $2r/k$ edges are added and each horizontal edge adds at most 4 vertices to the total. Since $r \geq r_0$ and $k \geq (r/2)^{2/3}$ we have $k \geq (r_0/2)^{2/3}$ and hence $r/k \leq r/(r_0/2)^{2/3}$. Given this bound on r/k , it is easy to see that it suffices to show that there is a positive constant α such that at least αr vertices lie on the boundaries of good chunks. As noted before, the dual graph of the chunks is a tree with at least r/k vertices. By the definition of special point, it is easy to see that chunks of degree 1 have at least $k/36$ vertices on their boundary. In addition, it is straightforward to check that, for any $m \geq 2$, the total number of vertices on the boundaries of a set of m chunks of degree 2 that form a path is at least $mk/36$. Putting these facts together with the easily proved lemma that, for any p -vertex tree T , the number of leaves in T plus the number of degree 2 vertices in T that are adjacent to a vertex of degree 2 is at least $p/4 + 1$, we see that at least $(r/4k)(k/36) = r/144$ vertices lie on the boundaries of good chunks. ■

3. The Splitting Algorithm

In this section we use the standard definition of horizontal visibility partition, which is trivially translatable to and from the two-sided wrap-around version in linear time. If E is a collection of edges that are disjoint except possibly at their endpoints, the **HVP** (horizontal visibility partition) is the partition obtained by adding horizontal edges as follows. For each endpoint z of an edge e in E , we add a horizontal line segment through z that extends in both directions until it intersects another edge of E on each side. We also add two dummy vertical edges at $x = -\infty$ and $x = +\infty$, so that each region (except for the top and bottom) in the HVP is a trapezoid. An example is shown in Figure 5. We will refer to the two non-horizontal edges of a trapezoid as the **sides** of the trapezoid, and will use the term **slice** an edge to mean dividing an edge into two edges by inserting a new vertex on the edge.




Figure 5. A horizontal visibility partition




Figure 6. The edge e is trapped by A .

For the horizontal visibility partitioning algorithm in the preceding section, we need an algorithm that, given the HVP of a polygonal chain C , and a subchain C_1 , computes the HVP of C_1 in time linear in the length of C . Let us call this the **chain splitting problem**. Our splitting algorithm actually works, and seems much easier to prove correct, in a more general context. To give this context, we must introduce the concept of hole-free. We will say that a set of edges A trap an edge e in the HVP of E if there is a simple polygon P containing e in its interior, where each edge in P is either horizontal and lies in the interior of a trapezoid of the HVP of E with both sides in A or is a subsegment of an edge in A . An example is shown in Figure 6. We say A is **hole-free with respect to E** if no edge of $E \setminus A$ is trapped by A . It is straightforward to check that modifying the set E (and possibly A) by either removing an edge of A or slicing an edge in two does not affect the condition that A is hole-free with respect to E . We define the degree of an edge e in E , denoted by $d(e)$, to be the number of trapezoids adjacent to it, and set $D(A) = \sum_{e \in A} d(e)$.

Suppose $E = E_1 \cup E_2$ is a collection of edges that are disjoint except possibly at their endpoints, such that E_2 is hole-free with respect to E . Our splitting algorithm solves the following problem which we will call **hole-free edge removal**. Given the HVP of E , compute the HVP of E_1 in time linear in the number of trapezoids adjacent to edges in E_2 .

It is worth observing that some restriction such as the hole-free condition must be made on the set of edges E_2 to allow their removal in linear time, since it is easy to give an $\Omega(n \log n)$ lower bound for the general problem of removing $n/2$ edges from an HVP of n edges by a reduction from sorting.

We begin by showing that chain-splitting is reducible to hole-free edge removal. Let C be a polygonal chain such that we know its HVP, and let C_1 be a subchain of C . Let E_1 be the set of edges in C_1 together with an infinite vertical edge L through the leftmost vertex on C_1 , let E_2 be the set of edges of C which are not on C_1 , with each edge split into two edges if it intersects L , and let $E = E_1 \cup E_2$. It is easy to see that we can obtain the HVP of E from the HVP of C , and the HVP of C_1 from the HVP of E_1 , in linear time. Moreover, because of the addition of L and the fact that $C_1 \cup L$ is connected, it is easy to see that no edge in E_1 is trapped by E_2 .

We now describe the hole-free edge removal algorithm. We use a simple data-structure to represent the HVP of E with the property that any edge can be removed in time proportional to its degree. Specifically, we represent each trapezoid, T , in the HVP by the list of its corner points in clockwise order, plus pointers to the edges to which it is adjacent. In addition, we add four pointers, left-up, left-down, right-up, and right-down. If e is the left side edge of T , the pointer left-up points to the trapezoid which also has e as its left side edge and is immediately above T in the HVP. The other pointers are defined analogously. Finally, for each edge e we provide pointers to its top and bottom trapezoids on both sides. Thus, in effect we have, for each edge e in E , two doubly linked lists of the trapezoids on its left and right sides, ordered according to their order along e . To remove an edge, e , we simultaneously walk along its two lists of adjacent trapezoids, extending each horizontal edge which ended at the interior of e through the trapezoid on the other side, splitting and merging the trapezoids from the two lists as necessary (see Figure 7). Since a bounded amount of work is needed for each horizontal edge with an endpoint on e , the total time to remove e is clearly $O(d(e))$.




Figure 7. The process of removing an edge

As a consequence of this data structure, for d a constant, we may assume that all the edges in E_2 have degree at least d , by pulling out small degree edges one at a time until no more remain. Our goal is to slice up the edges in E_2 so that if F is the set of pieces that share a trapezoid with an edge in E_1 , and E'_2 the pieces that do not, there is a fixed $\gamma < 1$ such that $D(E'_2) \leq \gamma D(E_2)$. Moreover, the slices will be chosen so that we can compute the HVP of $E_1 \cup F \cup E'_2$ from the HVP of E , and also the HVP of E_1 from the HVP of $E_1 \cup F$, in $O(D(E_2))$ time. It is not hard to prove that, by our previous observations on the preservation of the hole-free property under edge removal and slicing, and because every edge in F shares a trapezoid with an edge in E_1 , E'_2 is hole-free with respect to $E_1 \cup F \cup E'_2$. Thus we can recursively apply the hole-free edge removal algorithm to obtain the HVP of $E_1 \cup F$. It is easy to see that, so long as we achieve our slicing goal in $O(D(E_2))$ time, the algorithm runs in $O(D(E_2))$ time as desired.

In order to describe how the slicing of the edges in E_2 is done, we need to colour the trapezoids in the HVP of E . A trapezoid is white if at least one of its side edges is a subsegment of an edge in E_1 and grey otherwise (i.e. if both its side edges are subsegments of edges in E_2). See Figure 8(a) for an illustration. In the next two lemmas, using the hole-free property and the fact that every edge in E_2 has degree at least d , we prove an upper bound on the number of grey trapezoids.

Lemma 3.1. If $d \geq 9$, then every pair of edges in E_2 form the sides of at most one trapezoid in the HVP of E .

Proof. Suppose e, e' are edges in E_2 and h, h' are horizontal edges in the HVP of E such that both h and h' join e to e' with h the highest such horizontal edge and h' the lowest such horizontal edge. Let T be the trapezoid formed by h, h' and the appropriate subsegments of e and e' . It suffices to show that no edges of the HVP of E lie inside the interior of T . Since E_2 is hole-free with respect to E , no edges of E_1 lie inside T . Let m be the number of edges of E_2 lying inside T . We will show that $m = 0$. It is easy to see that the number of trapezoids of the HVP of E which lie inside T is at most $4m + 1$ because the m edges of E lying inside T generate at most $4m$ horizontal edges. This implies that if $m > 0$ then some edge inside T has degree at most 8, but this contradicts the assumption that every edge in E_2 has degree at least 9. ■

Lemma 3.2. If $d \geq 9$, then the number of grey trapezoids is at most $4D(E_2)/d$.

Proof. By the preceding lemma, the number of grey trapezoids is at most the number of trapezoids in the HVP of E_2 , which is at most $4|E_2|$. Finally, since each edge in E_2 had degree at least d we have $D(E_2) \geq d|E_2|$. ■

For each edge in E_2 , we slice it at each of its adjacent horizontal edges which border both grey and white trapezoids, to obtain the set of edges E_3 , and compute the new HVP of $E' = E_1 \cup E_3$ (see Figure 8(b)). Our data-structure supports doing this in time proportional to the degree of the edge in the current HVP, and it is not hard to check that the total amount of time used is proportional to $D(E_2)$. Let F be the set of edges in E_3 which share a trapezoid of the HVP of E' with an edge in E_1 , and let $E'_2 = E_3 \setminus F$. Note that because of the slicing used to form E_3 , every edge in F has at most one side on which it is adjacent to edges in E_3 .

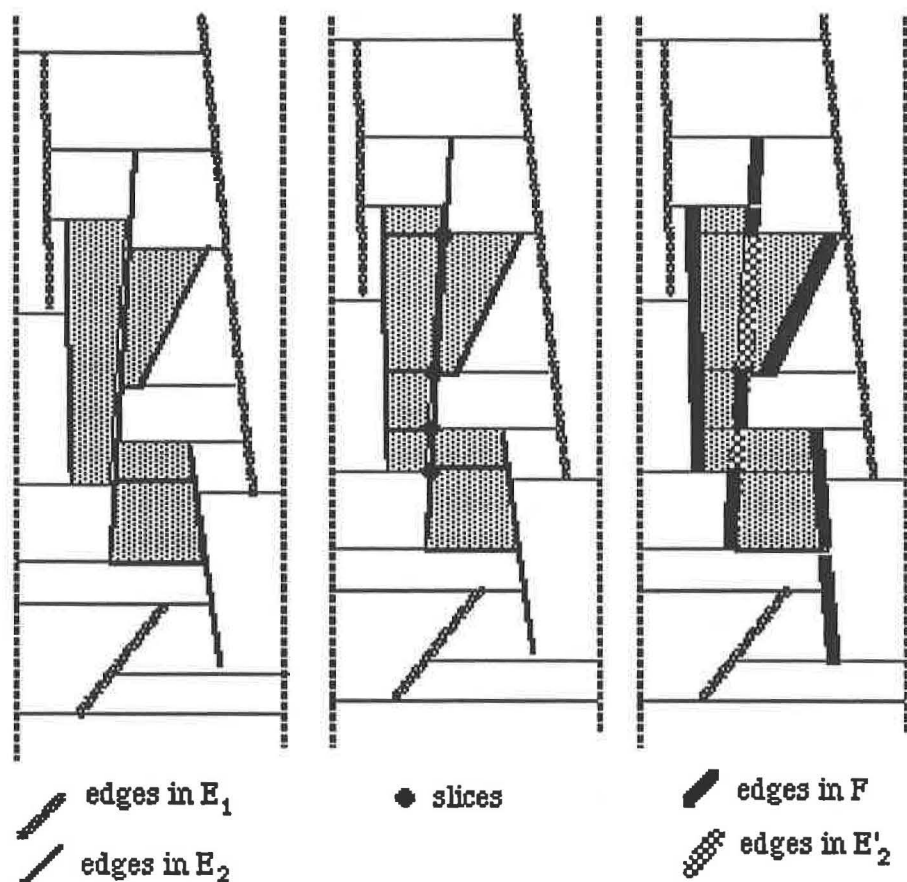


Figure 8.

Lemma 3.3. For $d \geq 25$ we have $D(E'_2) \leq 24D(E_2)/25$.

Proof. We colour a trapezoid in the HVP of E' black if it is adjacent to an edge in E'_2 . Let r [r'] be the number of grey [black] trapezoids in the HVP of $E[E']$. Since $D(E'_2) \leq 2r'$, by lemma 3.2 it suffices to show that $r' \leq 3r$.

We define the function q from the black trapezoids in the HVP of E' to the grey trapezoids in the HVP

of E as follows. Let U be a black trapezoid. If U is also a grey trapezoid then we set $q(U) = U$. If not, then at least one of the horizontal edges of U must be the extension of a horizontal edge h of the HVP of E that bordered both grey and white trapezoids, and we set $q(U)$ to be the grey trapezoid bordered by h . It is easy to see that for any grey trapezoid T , there are at most three black trapezoids that q maps to T . To be precise, for each horizontal edge h of T there is at most one black trapezoid with a horizontal edge formed by extending h , and adding the possibility that q also maps T to itself yields a maximum of three. ■

Finally we must show that the HVP of E_1 can be computed from the HVP of $E_1 \cup F$ in $O(D(E_2))$ time. Let f be the number of trapezoids adjacent to edges in F in the HVP of $E_1 \cup F$. It is straightforward to check that $f = O(D(E_2))$. Note that every edge in F now has at most one side on which it is adjacent to trapezoids adjacent to other edges in F . Thus we can partition $F = F_L \cup F_R$ according to whether this side is the left or right side respectively. Let f_L, f_R be the number of trapezoids adjacent to edges in F_L, F_R respectively in the HVP of $E_1 \cup F$. We will pull out all the edges in F one by one, doing all edges in F_L first. Removing an edge in F_L can only increase degrees of edges in $E_1 \cup F_R$, so the total time for removing the edges in F_L is $O(f_L)$, and the total increase of the degrees of edges in F_R is also $O(f_L)$. Now removing any edge in F_R only increases the degrees of edges in E_1 , so the total time for removing the edges in F_L is $O(f_L + f_R) = O(f)$. ■

4. An $O(n \log^* n)$ -Time Algorithm for Rasterized Polygons

We will call an n -vertex chain or polygon **rasterized** if the coordinates of its vertices are integers whose size is bounded by a fixed polynomial $p(n)$. In this section we show how our algorithm can be modified to compute the horizontal visibility partition of an n -vertex rasterized polygonal chain in $O(n \log^* n)$ time, where the constant depends on the polynomial $p(n)$.

The bottleneck of the algorithm given in section 2 is the finding of the horizontal neighbours of the special points (Lemma 2.2). More precisely, this is the only place where we are forced to have the length of segments, to which we apply the algorithm recursively, be large, in order to have sufficiently few special points so that we can find their horizontal neighbours in linear time. The key to obtaining an $O(n \log^* n)$ time algorithm for rasterized polygons is that given any n -vertex rasterized chain we can construct a data-structure in linear time such that the horizontal neighbours of any point can be found in $O(\log^2 n)$ time. This will allow us to make the length of the segments be $\log^2 n$ instead of $n^{2/3}$, and still find the horizontal neighbours of the $2n/\log^2 n$ special points in $O(n)$ time.

We will say that a rational number is $p(n)$ - **bounded** if the absolute values of its numerator and denominator are both bounded by $p(n)$. The key property of a rasterized chain or polygon is that the x -coordinate of the intersection of any edge with the horizontal line $y = y_0$ where y_0 is a $p(n)$ -bounded integer, will be a $(p(n))^2$ -bounded rational. Moreover, we can sort n $(p(n))^2$ -bounded rationals in $O(n)$ time (see Lemma 3.2 in [KK81] for example). We will say that a set of edges is **inner-disjoint** if they are disjoint except possibly at their endpoints. We will use the term **fence** for a set of inner-disjoint edges $\{e_1, \dots, e_k\}$ and a horizontal line L that intersects all the e_i . We will refer to the line L as the **wire** of the fence, and the x -coordinate of the intersection of an edge of the fence with the wire as its **wire-coordinate** . We will say the fence is

ordered if we know both the sorted order of the vertical coordinates of the endpoints of the e_i , and the left to right order of the wire coordinates of the e_i .

Lemma 4.1. If $F = (\{e_1, \dots, e_k\}, L)$ is an ordered fence, then in $O(k)$ time we can construct a data structure such that for any point z we can find its horizontal neighbours in F in $O(\log k)$ time.

Proof. It suffices to show how to obtain the HVP of F in $O(k)$ time since given the HVP we can use planar point location [EGS86] to construct a data-structure to answer queries in the HVP of F in $O(\log k)$ time. By symmetry it suffices to find the horizontal neighbours of the top endpoints of the edges in F . Since we know both orders, in $O(k)$ time we can create a data-structure with a doubly linked list joining the edges in left to right order according to their wire-coordinates, and a linked list connecting the edges according to increasing vertical coordinate of their top endpoint. We now obtain the left and right neighbours of each top endpoint in increasing order of vertical coordinate, by assigning the current lowest top endpoint its left and right neighbours in the doubly linked list as its horizontal neighbours, and then deleting that endpoint's edge from the data-structure. Since this can be done in constant time for each endpoint the total time used is $O(k)$. ■

It is possible to prove the preceding lemma without the assumption that we know the sorted order of the vertical endpoints, by reducing finding the HVP of F to finding the HVP of a monotone polygon. However, since we can easily obtain the sorted order of the vertical coordinates we use this version since the proof is so easy and direct.

Lemma 4.2. Suppose E is a set of inner-disjoint edges such that we know the sorted order of the vertical coordinates of the endpoints of the edges. Then in $O(|E|)$ time we can partition E into a family of fences, $F^{i,j}$, such that for any point z there are $O(\log |E|)$ fences in the family which can contain the horizontal neighbours of z , and those fences can be determined in $O(\log |E|)$ time.

Proof. Let $m = |E|$ and let y_1, \dots, y_{2m} be the vertical coordinates of the endpoints of the edges of E in (increasing) sorted order. We can assume we know the rank of each vertical coordinate, since otherwise we can compute this in $O(|E|)$ time from the sorted order. For each edge e let $l(e) = e_2 - e_1$, where y_{e_1} and y_{e_2} are the vertical coordinates of the two endpoints, with $y_{e_1} \leq y_{e_2}$. We define $F^{i,j} = \{e : 2^i \leq l(e) < 2^{i+1}, j2^i \leq e_2 < (j+1)2^i\}$. It is easy to see that each non-empty $F^{i,j}$ is a fence since every edge intersects the line $y = y_{j2^i}$. For each edge in E we can clearly determine to which $F^{i,j}$ it belongs in constant time, and hence we can obtain the partition into fences in $O(|E|)$ time. Finally, for any point z let z_y be the vertical coordinate of z . By binary search we can determine the index p such that $y_p \leq z_y < y_{p+1}$ in $O(\log |E|)$ time. Now for each i there are at most two values of j such that $F^{i,j}$ can contain a horizontal neighbour of z , and these can clearly be determined in constant time once we know p . ■

The construction used to separate the edges into fences in the preceding proof is essentially equivalent to (static) interval trees, which were introduced by Edelsbrunner [E80].

Lemma 4.3. Suppose E is a set of inner-disjoint rasterized edges, and that we are given a partition of E

into disjoint sets, $E = E_1 \cup \dots \cup E_t$. Specifically, for each edge e in E we are given the index i of the set E_i containing e . Then in $O(|E|)$ time we can construct data structures D_1, \dots, D_t such that for any point z we can find its horizontal neighbours in E_i in $O((\log |E_i|)^2)$ time.

Proof. First note that in $O(|E|)$ total time we can obtain, for each E_i , the order of the vertical coordinates of its edges, by performing a double radix sort. Next, using the technique described in Lemma 4.2, we partition the edges in each E_i into the appropriate fences. Now for each edge in E we compute its wire-coordinate with respect to its fence. Again by performing a double radix sort, in $O(|E|)$ time we compute the sorted order of wire-coordinates and the sorted order of the vertical coordinates of the endpoints of the edges in each fence. Next, by Lemma 4.1, in $O(|E|)$ total time, we construct a data structure for each fence such that for any point we can find its horizontal neighbours in the fence in time logarithmic in the number of edges in the fence. Now for any point z and any i , we can find the horizontal neighbours of z in E_i as follows. In $O(\log |E_i|)$ time we identify the $O(\log |E_i|)$ fences of edges in E_i which could possibly contain the horizontal neighbours of z . Now, for each of these fences, in $O(\log |E_i|)$ time we locate the horizontal neighbours of z in the fence. This takes a total of $O((\log |E_i|)^2)$ time. Finally we compare the horizontal neighbours found in each of these fences to find the true horizontal neighbours of z in E_i . ■

Given Lemma 4.3, it is fairly straightforward to modify the algorithm in Section 2 to run in $O(n \log^* n)$ time on rasterized chains. The overall structure of the algorithm is that we partition the chain into subchains of length $\log^2 n$, apply the algorithm recursively to compute the WP of each subchain, and then compute the WP of the entire chain using a modification of Theorem 2.1, which we state below.

Theorem 2.1 (rasterized version). Suppose $k \geq \log^2(r/2)$. Given a k -uniform partition of an r -vertex non-degenerate polygon P and a data-structure so that for any point z we can find the interior horizontal neighbours of z in P in $O(k)$ time, we can compute the WP of the interior of P in $O(r)$ time.

The proof of this version of Theorem 2.1 is identical to the proof of the previous one except for two points. First, the assumption about the data-structure eliminates the need for Lemma 2.2, and second we need to have data structures for the bad chunks in order to apply the theorem recursively. However, this is trivial since the data-structure for P works for all the chunks as well. Thus, the only remaining issue is how to ensure that the necessary data-structure is available every time Theorem 2.1 is applied at a “top” level. We do this in $O(n \log^* n)$ time by $\log^* n$ applications of Lemma 4.3 right at the beginning of the overall algorithm, so that all the necessary data structures are computed before doing anything else. More precisely, we first let E be the set of edges in the doubling of the original chain and apply Lemma 4.3 to the trivial partition $E = E_1$. This provides the data-structure that will be needed when Theorem 2.1 is applied to the doubling of the original chain.

We next partition the chain into subchains of length $\log^2 n$ and let $E = E_1 \cup \dots \cup E_t$ where E_i is the set of edges in the doubling of the i -th subchain. Applying Lemma 4.3 provides the necessary data structures for the applications of Theorem 2.1 where the polygon is the doubling of one of these subchains. We repeat this process with partitions of the original chain into subchains of length $\log^2(\log^2 n)$, then $\log^2(\log^2(\log^2 n))$, etc., corresponding to the recursive structure of the overall algorithm. Since each application of Lemma 4.3 requires $O(n)$ time, the total time needed is $O(n \log^* n)$.

REFERENCES

- [BT86] Bhattacharya, B.K. and Toussaint, G.T., "A linear algorithm for determining the translation separability of two simple polygons", McGill University, School Comput. Sci., Report SOCS-86.1, Montreal, Canada, 1986.
- [C82] Chazelle, B., "A theorem on polygon cutting with applications", *Proc. 23rd Annual Symp. on Found. of Comput. Sci.*, 1982, pp. 339-349.
- [C184] Chazelle, B. and Incerpi, J., "Triangulation and shape complexity", *ACM Trans. on Graphics*, 3 (1984), pp. 135-152.
- [C90] Chazelle, B., "Triangulating a simple polygon in linear time", *Princeton Tech. Report CS-TR-264-90*, 1990.
- [CTV88] Clarkson, K., Tarjan, R. and Van Wyk, C., "A fast Las Vegas Algorithm for triangulating a simple polygon", *Proc. 4th ACM Symp. Comp. Geom.* (1988) pp.18-22, *Princeton Tech. Report CS-TR-157-88*, 1988.
- [EGS86] Edelsbrunner, H., Guibas, L.J. and Stolfi, J., "Optimal point location in a monotone subdivision", *SIAM J. Comput.*, 15 (1986), pp. 317-340.
- [FFR83] Fiume, E., Fournier, A. and Rudolph, L., "A parallel scan conversion algorithm with anti-aliasing for a general-purpose ultra-computer", *ACM Comput. Graphics*, 17, 3 (July 1983), pp. 141-150.
- [FM84] Fournier, A. and Montuno, D.Y., "Triangulating simple polygons and equivalent problems", *ACM Trans. on Graphics*, 3(1984), pp. 153-174.
- [FNTV88] K.Y. Fung, T.M. Nicholl, R.E. Tarjan, and C.J. Van Wyk, "Simplified linear-time Jordan sorting and polygon clipping", *Princeton Tech. Report CS-TR-189-88*, 1988.
- [GHLST87] Guibas, L., Hershberger, J., Leven, D., Sharir, M. and Tarjan, R.E., "Linear time algorithms for visibility and shortest path problems inside triangulated simple polygons", *Algorithmica* 2 (1987), pp. 209-233.
- [GJPT78] Garey, M.R., Johnson, D.S., Preparata, F.P. and Tarjan, R.E., "Triangulating a simple polygon", *Inform. Process. Lett.*, 7 (1978), pp. 175-180.
- [E80] H. Edelsbrunner, Dynamic data structures for orthogonal intersection queries, Techn. Univ. Graz, Instit. Informationsverarb., Report f59, Graz, Austria, 1980.
- [HM83] Hertel, S. and Mehlhorn, K., "Fast triangulation of a simple polygon", *Proc. Conf. Found. of Comput. Theory*, New York, *Lecture notes on Computer Science* 158 (1983), pp. 207-218.
- [HMRT86] Hoffman, K., Mehlhorn, K., Rosenstiehl P., and Tarjan R., "Sorting Jordan sequences in linear time using level-linked search trees", *Inform. and Control* 68(1986), pp.170-184.
- [KK81] J.M. Keil and D.G. Kirkpatrick, "Computational Geometry on Integer Grids", *Proc. 19th Annual Allerton Conference on Communication Control and Computing*, September, 1981, pp. 41-50.
- [K83] Kirkpatrick, D.G., "Optimal search in planar subdivisions", *SIAM J. Comput.* 12, 1 (Feb. 1983), pp. 28-35.
- [L81] Lee, D.T., "Shading of regions on vector display devices", *ACM Comput. Graph.*, 15, 3 (July 1981), pp. 34-44.
- [LTL89] Liou, W.T., Tan, J.J.M. and Lee, R.C.T., "Minimum partitioning simple rectilinear polygons in $O(n \log \log n)$ time", *Proc. 5th ACM Symp. Comp. Geom.* (1989) pp.344-353.
- [TV88] Tarjan, R.E. and Van Wyk, C.J., "An $O(n \log \log n)$ -time algorithm for triangulating a simple polygon", *SIAM J. Comput.* 17, 1 (Feb. 1988), pp. 143-178.
- [T80] Toussaint, G., "Pattern recognition and geometrical complexity", *Proc. 5th Intern. Conf. Pattern Recognition* 1980, pp.1324-1347.
- [T86] Toussaint, G., "Computational geometry and morphology", *McGill Univ. Tech. Report SOCS-86.3*, Montreal, Canada 1986.
- [T88] Toussaint, G., "An output-complexity-sensitive polygon triangulation algorithm", *McGill Univ. Tech. Report SOCS-88.11*, Montreal, Canada 1988.
- [T88₂] Toussaint, G., (Ed.), *Computational Morphology*, North-Holland, 1988.
- [TA82] Toussaint, G. and Avis, D., "On a convex hull algorithm for polygons and its application to triangulation problems", *Pattern Recognition* 15,1(1982), pp.23-29.

